

IV.B.1.f QUANTUM CHEMICAL INVESTIGATIONS

QUANTUM CHEMICAL INVESTIGATIONS OF
HEME PROTEINS AND FERREDOXINS

Dr. Gilda Loew
Stanford Department of Genetics

(Grant NSF GB-40105, 2 years, \$18,000 this year)

SUMEX is used for the calculation of various one-electron electromagnetic properties of iron containing compounds. The programs were formulated and written by David Steinberg, Michael Chadwick, and David Lo. David Lo was responsible for converting the programs for interactive use on the PDP system. Slight improvements were made by Robert Kirchner, and Sheldon Aronowitz is currently expanding the formulation to include additional spin and oxidation states of the iron atom.

The properties that are calculated include the electric field gradient at the iron nucleus, quadrupole splitting, isotropic and anisotropic hyperfine interaction, spin-orbit coupling and zero field splitting, g values, and temperature dependent effective magnetic moments. The calculated values are compared directly to experimental results obtained from published Mossbauer resonance and electron spin resonance spectra. Such a comparison determines not only the reliability with which these properties can be calculated but also gives an indication of the ability of the model of the iron active site to mimic the actual environment found in a particular compound or iron containing protein.

The major input to these properties programs is a description of the electron distribution of the compound under consideration. This description is obtained using a semi-empirical molecular orbital method employing the interactive extended Huckel procedure. Such a calculation requires up to 660K core and is performed elsewhere. When the calculated electron distribution yields a set of calculated properties in agreement with observation, we have increased faith in the description of the model of the active site and can carry the model one step further to make qualitative inferences about certain properties relevant to the biological functioning of the compound. These properties, which are harder to characterize experimentally, include the nature of the ligand binding to iron, relative bond strengths themselves, net atomic charges, and electric potentials. The model may be varied (that is, change the spin or oxidation state of the iron, replace certain ligands, or simply change the geometry of the ligands) and a new set of properties calculated to predict what effect these changes would have on the observed electromagnetic properties.

Such a procedure lends itself well to the study of three classes of iron containing compounds of biological interest: the one-iron sulfur proteins known as rubredoxins and the two-iron sulfur proteins called

plant-type ferredoxins which serve as one electron transfer agents, heme proteins which serve as oxygen as well as one electron transfer agents, and sideramines which serve as iron transport agents. The calculated properties for the first class, used to elucidate the geometry of the sulfur ligands and the spin state of the iron within the protein, are reported in the following

PUBLICATIONS:

- [1] G.H. Loew, M. Chadwick, and D.A. Steinberg, Theoret. Chim. Acta (Berl.) 33, 125 (1974)
- [2] G.H. Loew and D.Y. Lo, Theoret. Chim. Acta (Berl.) 33, 137 (1974)
- [3] G.H. Loew, M. Chadwick, and D.Y. Lo, Theoret. Chim. Acta (Berl.) 33, 147 (1974)
- [4] G.H. Loew and D.Y. Lo, Theoret. Chim. Acta (Berl.) 32, 217 (1974)

We are currently performing a systematic study of heme proteins. The electromagnetic properties of these proteins and of synthesized compounds which mimic the observed behavior of the proteins have been well studied experimentally. But many questions regarding the nature of small ligand binding to the heme group remain unresolved. Before we can address ourselves to such problems, we must first be able to theoretically reproduce the experimentally observed behavior. The specific areas of interest are:

- a.) deoxy heme in both the relaxed (iron in the plane with a low or high spin state) and tense (iron out of the plane with a high spin state) configurations.
- b.) oxy heme with various oxygen geometries (co-axial or coplanar) and several excited electronic states (promotion of an electron from an iron d-orbital to an unfilled oxygen orbital).
- c.) abnormal heme compounds which do not bind oxygen but which bind axially to CN, N₃, NO, or OH.
- d.) the enzymatic cycle of cytochrome P450 camphor, in which the protein has been isolated with the iron in various spin and oxidation states.

Preliminary results for oxy heme have been published (Loew and Kirchner) in the J. Amer. Chem. Soc. 97, 7388 (1975).

We have also calculated the electromagnetic properties of ferrocene Fe(C₅H₅)₂ and the binuclear transition metal complexes biferrocene and biferrocenylene in various oxidation states. The work has been published (Kirchner and Loew) in Theoret. Chim. Acta (Berl.) 41, 1 (1976) and submitted (Kirchner and Loew) to Inorganic Chemistry.

The heme work is funded by the National Science Foundation Grant GB 40105, which was renewed starting June 1976 for a period of two years.

Undergraduate research projects which attempt to correlate reactive electronic sites for a series of polycyclic aromatic hydrocarbons with carcinogenic activity use SUMEX to calculate various measures of C-C and C-H bond reactivity. Again, the major input to this program is taken from the results of an iterative extended Huckel molecular orbital calculation performed elsewhere. The only current funding for these projects is a SCIP computer processing subsidy although an application to the National Institute of Health is pending.

IV.B.2 NATIONAL PILOT PROJECTS

Over that past year several national pilot projects have been initiated with the approval of the AIM Executive Committee and advice of the AIM Advisory Group. One of these (the ACT Project under Dr. John Anderson) has moved from pilot status to become a formal project. The currently active pilot efforts are summarized below.

IV.B.2.a NATURAL LANGUAGE UNDERSTANDING

Natural Language Understanding

Prof. R. Lindsay
University of Michigan

(Financial support from University of Michigan)

I. Summary of Research Program

The major aims of this pilot project have been to establish research goals, initiate collaborations with faculty at the University of Michigan Medical School, and to develop software. The project staff consists of Associate Professor Robert K. Lindsay, Dr. Maija Kibens (Research Associate), and Mrs. Kathie Gourlay (Programmer Analyst), all of the University of Michigan.

A) Technical Goals

The overall goal of this project is the development of a histological model that will assist anatomists in the design and evaluation of methods of organ culture. As conceived at present, our technical goals are the design of (a) a data structure for encoding descriptions of microscope slides made from organ explants, (b) a model of microanatomical processes based on the expertise of histologists and pathologists, and (c) means for construction of the data structures of (a) from histologists' verbal descriptions.

B) Medical Relevance and Collaboration

The value of such a system to biology and medicine is far-reaching to the extent that it succeeds in assisting in the development of organ culture methodology. To illustrate with a single important example, the ability to cultivate the organs of experimental animals is the first step in the in vitro study of disease processes such as cancer in those organs. We are working in collaboration with three anatomists in the Department of Anatomy, Professor Raymond Kahn, Associate Professor William Burkel, and Assistant Professor Theodore Fischer. For the past two years this group has been experimenting with methods for cultivating canine prostate.

C) Progress and Accomplishments

Our efforts have been directed along two fronts. We are familiarizing ourselves with the current capabilities, knowledge, and problems of the histology group. We are also developing artificial intelligence programs to understand histological information typed in a natural language format.

Collaborating with the histologists

1. We are interviewing the principal investigators and their several assistants individually to learn from each of them his conception of tissue functioning.
2. Formulating better analysis methods - Together with the histologists we have designed a new grading scheme for recording the maintenance status of an organ explant. This scheme is more complex than their previous category system in that it includes more factors and finer distinctions. Currently, the group is using standard non-parametric statistics to compare the evaluations of the explants.
3. Designing an AI model - The histologists are enthusiastic about the new evaluation method. They believe it to be a great improvement over their previous one. However, they recognize its inadequacies and are open to any artificial intelligence techniques that will enable them to capture more of their knowledge about each explant in a form that can be used in the design and evaluation of experiments.

Software Development for Natural Language Input

1. Formulating a general design - The proposed structure for the system includes multiple sources of knowledge, each contributing hypotheses about the meaning of the input. The input is to be freely formatted with possible typing, spelling, and syntactic errors. The output will be an internal representation of the meaning of the input text, namely a representation of an organ explant. The knowledge will include components for: typing correction, word decomposition, morph recognition, syntactic analysis, semantic analysis, and histological knowledge. The knowledge components of the system are being designed to be independent of each other insofar as possible.
2. Implementing - The knowledge components for typing correction and word decomposition have been written in INTERLISP. The dictionary format has been designed. Work is currently being done on the morph recognition component.

D) Publications

A manuscript describing the results of the application of the revised evaluation scheme has been written by Professors Kahn, Burkel, Fischer, and Herwig (the project surgeon). The paper is titled "Effect of

vitamin A on canine prostate in organ culture". There have been no publications to date on the AI aspect of this project.

E) Funding Status

The canine prostate project is funded by the National Cancer Institute. A grant application to NIH for similar work with lung is pending. The salaries and research facilities of Professor Lindsay, Dr. Kibens, and Mrs. Gourlay are provided by the University of Michigan Mental Health Research Institute, a division of the Department of Psychiatry in the Medical School.

II. Interactions with the SUMEX-AIM resource

A) Medical Use of Programs through Networks

We have not had occasion for such use.

B) Useful Contacts and Cross Fertilization with Other SUMEX-AIM Projects

Kathie Gourlay is on the LISP users' mailing list. Most of the other users and personnel at SUMEX have been very helpful in giving advice and solving problems. A few examples of such assistance during the past year follow.

There have been some problems with the speed of response in an INTERLISP program. Masinter has been very helpful with suggestions. In one instance, he was able to run the program and interactively pinpoint some slow spots.

We have also had communication via SNDMSG with N. Smith, Hedberg, Feigenbaum, Davis, Lederberg, R. Smith, Colby, Parkison, Winograd, and others.

Dr. Kibens has used the LINK facility to obtain answers to questions about details of system use, and has used SNDMSG extensively for such purposes. Other interactions have concerned obtaining information about current status of natural language input systems for interactive programs. SUMEX has also been a very convenient communications facility via SNDMSG to non-SUMEX users at SRI, CMU, and SU-AI.

C) Critique of Resource Services

The SUMEX staff has been very helpful. To cite just one example: At one time last year, we wanted to transfer some data to SUMEX from a PDP-9 minicomputer which is located here. Gourlay communicated with Cower via SNDMSG and arranged to mail a DECtape. Contrary to what the DECsystem 10 Assembly Language Handbook led Cower to believe, the PDP-9 DECtape was not compatible with the PDP-10 software. Cower and another person spent

several hours working unsuccessfully to transfer the data. We appreciate their efforts although the problem could not be solved because of serious incompatibilities.

Now that Tymnet has several local lines to our area (since January 1976) and we have terminals located in our offices the SUMEX facility is very convenient. The system is quite reliable. However, when it is down the explanation given us from Tymnet is almost always out-of-date, e. g., maintenance work that was completed hours ago.

There should be a manual such as the Tenex Executive Manual that explains the features available on the SUMEX system. At present, it is necessary to look in dozens of different documentation files or to learn by hearsay. In the interim, it would be good to have one of the system personnel designated as a general source person for details of system use.

We would suggest that some thought be given to making the LINKing facility a more productive and convenient means of communication. While LINKing is a potentially useful device, it is also a potential nuisance to the recipients. This is a cause of our reluctance to use this facility. Perhaps an explicit policy should be decided upon by all SUMEX users to establish what is the community attitude toward LINKing. It might facilitate communication by LINKing if the default option were changed to REFUSE, but modified to allow immediate acceptance of the LINK upon learning who it is from. Certain programs, such as TYPE, are usually run in REFUSE mode. Perhaps these programs should set REFUSE mode on entry and clear it on exit so that the user would be protected from interruption at those times without needing to have the REFUSE mode set permanently. There are obviously many such technical changes that could be made to improve this feature, and we would like to see some discussion of them.

We look forward to the availability of 1200 baud capability over TYMNET so that listings can be obtained more rapidly. Any encouragement from SUMEX to TYMNET that would speed the conversion would be appreciated. We think it would be wise for the system to be compatible with the VADIC protocol (soon to be adopted by Bell, we hear unofficially) rather than with the troublesome Bell 202 equipment, as announced.

IV.B.2.b KRL PROJECT

Knowledge Representation Language - KRL

Dr. Dan Bobrow, Xerox PARC
Dr. Terry Winograd, SU AI Lab

This pilot project was just initiated on the SUMEX-AIM facility and is a medicine-oriented extension of the KRL development effort at Xerox Palo Alto Research Center. The basis of the original project is the development of a systematic programming framework within which to describe and manipulate knowledge about a task domain and which may be used by a performance program to reason and solve problems within that domain. The development of such AI tools is an important part of the AIM community in that it allows the more coherent and general formulation of medical AI programs.

A first version of KRL has been implemented and several students will experiment with implementing medical consultation programs (e.g., MYCIN, CASNET, or Rubin's model of renal disease) using KRL.

IV.B.2.c COMPUTERIZED PATIENT MONITORING

Computerized Patient Monitoring and Clinical Decision Making

John J. Osborne, M.D. Director Intensive Care Services
Richard R. Mitchell, Ph.D. Biomedical Engineer

The Institutes of Medical Sciences
Pacific Medical Center

The immediate desire of this pilot project is to use MLAB and to explore the opportunity to become a regular member of the SUMEX-AIM user community.

The project is part of a Bioengineering and Computer Science Resource for medical research. The Research Data Facility of the Institutes of Medical Science is a NIH funded center for the development of computerized patient monitoring and clinical decision making. The emphasis to date has been in the area of clinical monitoring of the respiratory parameters of critically ill patients.

There are three major areas of potential joint cooperation with SUMEX-AIM:

1. Clinical Decision Making;
2. Image Processing;
3. Modeling.

In the area of Clinical Decision Making the project is funded to develop an intelligent system for detecting and reporting alarm conditions in the hospital intensive care environment. Its goals include using sophisticated image processing techniques for the evaluation of pulmonary physiology using the scintillation camera and ¹³³Xenon. The present work in modeling is limited by the capabilities of the IBM 1800 CSMP and the investigators are interested in exploring the use of more complicated models requiring a sophisticated simulation language.

IV.B.2.d AI IN PSYCHOPHARMACOLOGY

Artificial Intelligence in Psychopharmacology
(NIH grant application in preparation)

Dr. Jon F. Heiser, M.D.
Assistant Adjunct Professor
Dept. of Psychiatry and Human Behavior
University of California at Irvine

A. Introduction

This project has just been authorized as an AIM project. The following quote from a letter of Drs. Buchanan and Axline of the MYCIN project describes the collaboration that lead to their project and which is expected to continue.

"He is extending MYCIN's knowledge base to cover consultations regarding chemotherapy for psychiatric disorders. This is valuable to us for at least two reasons: it increases the potential uses of the program and it illuminates those specific parts of the program that are not yet general enough to be easily extended to new areas. By pioneering in the effort to develop a more general framework for medical reasoning computer programs, Dr. Heiser is helping us provide a means for encoding and testing large amounts of medical knowledge."

The objective of the new project is to develop computer based automated systems capable of assisting in research, teaching and consultation in psychopharmacology. It will result in the development of software which will run on the University of California, Irvine PDP-10.

[The following material is abstracted from Dr. Heiser's proposal to the AIM Executive Committee].

A.1 BACKGROUND:

Information in medicine expands so rapidly that both researchers and clinicians struggle to digest it and apply it wisely. Computer-based instruction (textbooks, journals, individualized supervision and consultation) is one solution to this problem. By their very nature computer-based systems can be programmed to 1) explain their reasoning in natural language and in terms intuitively acceptable to users of various degrees of sophistication, 2) have their behavior totally analyzed, and 3) be easily modified or updated.

Computer-based knowledge systems have been developed for describing and solving problems in pharmacology. At Stanford University Medical Center in California, when new drugs are prescribed for a patient, their profile of action is compared to that of drugs already consumed by the patient. A warning is generated for the physician if a potential interaction is noted (1). Artificially intelligent systems have been

developed which utilize pharmaco-kinetic models to suggest initial doses and monitor on-going maintenance doses with complex drugs such as digitalis (2). Artificial intelligence systems are also being generated to diagnose patients with infectious diseases and to suggest appropriate antibiotic agents (3).

Several other systems were discussed at the First Annual AIM (Artificial Intelligence in Medicine) Workshop, held at Rutgers University 14 June through 17 June 1975. (S. Amarel and C.A. Kulikowski of the Computer Science Department at Rutgers University directed the conference).

We have begun to adapt the techniques of the Stanford Group (3,4) in the generation of an artificially intelligent system which evaluates and diagnoses psychiatric patients, suggest pharmacological treatment and monitors the on-going clinical course.

The system has 16 rules, based on conventional clinical observations, for diagnosing either mania or schizophrenia. Clinical findings are collected and a diagnosis made by manipulating human expert generated "certainty factors" which are similar to but not identical to probabilities. their precise mathematical nature and manipulation are described in Shortliffe and Buchanan (4).

A.2 RATIONALE:

Psychopharmacological agents are frequently misused qualitatively and quantitatively by prescribing physicians as well as by consumers. Consultation with experts in psychopharmacology is frequently sought and is given on the basis of clinical data, currently established practice, evolving research or ad hoc hypotheses. A computer based consultation system, available 24 hours per day, could greatly assist non-specialist physicians in choosing the best psychopharmacological treatment, given the same expertise and data. Such a system could also serve as a teacher-advisor to students and as a reference for various types of psychopharmacological knowledge, e.g., well established principles and practice, new but not fully verified ideas and late breaking developments (5). Properly weighted, all such information could be used in consulting and teaching functions. For example, the system could suggest less well established, more controversial or more hazardous diagnostic or therapeutic techniques for a patient with a life-threatening situation not responsive to conventional measures. Like the human clinician, in desperate or excessively chronic circumstances, the system could generate novel hypothesis with an estimate of potential risks and benefits.

B. SPECIFIC AIMS:

1. To study existing automated systems, computer based and otherwise, which assist in clinical decision making.
2. To develop a model of expert clinical decision making for clinical psychopharmacology.

3. To implement this model on a computer system such that the system can converse in real time in natural language through computer terminals with users located close by or remotely.
4. To evaluate the performance of the system as a teaching and consulting aide.
5. To increase the breadth and depth of the artificially intelligent system by a) increasing the technical sophistication, e.g., by adding options such as voice activated microphone-loudspeaker (or print-out) terminals. b) adding other areas of clinical psychiatry towards an ultimate goal of having a fully automated and self-contained textbook-consultant for psychiatry. c) linking the system to precoded data bases so that the system could quickly learn from thousands of actual case histories and use this data and experience both to modify its intuitively human-like rule-based decision model and to generate abstract mathematical or statistical decision making models. d) integrating the system with biomedical data collecting techniques achieve more direct involvement in research and patient evaluation. e) proposing new drugs to be synthesized or tested for desired psychopharmacologic affects.

Aims 1-4 should easily be attained within five years. Aim number 5 is obviously quite speculative and dependent on developments in computer technology, biomedical engineering, etc. Promising beginnings have been made. An example is a program which reads and analyzes the content of typed scripts of spontaneous human speech (6). This system parses sentences into noun and verb clauses, recursively or repeatedly if necessary, and uses a set of rules to score the noun and verb phrases for a variety of affects and states of mind by means of a well documented, reliable and valid technique of content analysis otherwise requiring a human content analyzer with a common sense knowledge of the world and the language.

C. METHODS AND PROCEDURES:

We plan to study existing systems and to develop a similar system in clinical psychopharmacology. Dialogues, equivalent to those planned for this development, are routinely produced by the antibacterial program mentioned above in (3) and could be available in the area of clinical psychopharmacology within a year in complete enough form to be evaluated. This work will be done in consultation with the Information and Computer Science Department at University of Calif., Irvine and the Computer Science Department at Stanford University. During approximately the first two years efforts will be concentrated on developing the artificial intelligence techniques referred to above, including question answering in natural language, abstract reasoning, advice giving, etc. The expert information will be installed in sentence-like form initially from the working knowledge of the principal investigator and the behavior of the system compared to the behavior of the principal investigator. In the second phase expert information will be abstracted from standard textbooks, journals and consultation with acknowledged experts. Here evaluation becomes more difficult except when the system makes an obvious

mistake. A third level of expert input will include other and possibly non-human information sources such as actuarial formulas and other statistical techniques (12).

In later phases of the project more concern will be placed on the diagnosis problem. This problem is being deemphasized during initial phases because it has received reasonable attention by other groups (7-12).

Many of the above mentioned systems have been formally or informally evaluated and found to perform, within their range of applicability and our ability to measure performance, as well as acknowledged experts (12). Consultation with experts in evaluation research in both education and clinical medicine is available locally and will be enlisted in later aspects of the project once a workable system has been developed.

Risks and hazards in this procedure are minimal since no biological material is involved, no patient records are used and identification of patients to be discussed can be secured or eliminated with no impact on the system or the user. Use of a consulting and teaching system in clinical psychopharmacology might involve clinical responsibility and could be regarded as contributing to or responsible for an error of omission or commission in clinical judgement and practice. Every attempt will be made to complete a thorough evaluation of the system, its validity and reliability before it is made available to other than a small testing group. If in later phases we add the capability of utilizing large data bases such as available through the Missouri Information System (10,12) only those well developed procedures for transmitting large data bases with complete anonymity and protection of individual patient rights will be used. Such procedures will be rigidly and consistently adhered to in all aspects of this project.

D. SIGNIFICANCE:

It is hoped that results from these initial studies will stimulate further research by physicians and graduate students in related fields such as biochemistry, pharmacology, pharmacy, mathematics and the information processing sciences. To have instant access to teaching and consultation based on a consensus of the best data, the best abstract mathematical or statistical "number crunching" techniques and the best human experts would be of great value to researchers, specialists, family physicians, students and others. Evaluation of the effects of such a system and comparison with traditional methods of research, teaching and consultation would be of great benefit to medical educators. It is hoped that many students would master, beat or "psyche out" the system. This would be excellent evidence that learning is occurring. However, because of the information explosion, periodic updates to and from the system should prevent it from becoming obsolete.

References

1. Cohen, S.N. et al. Computer-based monitoring and reporting of drug

- interactions. Proceedings MEDINFO IFIP Conference, Stockholm, Sweden, August 1974.
2. Silverman, H. A digitalis therapy advisor. MAC TR-143, Massachusetts Institute of Technology, Cambridge, Massachusetts, January 1975.
 3. Shortliffe, E.H., AXLINE, S.G., Buchanan, B.G. and Cohen, S.N. Design considerations for a program to provide consultation in clinical therapeutics. Proceedings of the San Diego Biomedical Symposium, February 1974, 311-319.
 4. Shortliffe, E.H. and Buchanan, B.G. A model of inexact reasoning in medicine. Mathematical Biosciences 23, 351-379, 1975.
 5. Ayd, F.J. Rules for neuroleptic therapy. International Drug Therapy Newsletter 9, 33-35 (1974).
 6. Gottschalk, L.A., Hausmann, C. and Brown, J.S. A computerized scoring system for use with content analysis scales. Comprehensive Psychiatry 16, 77-90, 1975.
 7. Johnson, J.H., Giannetti, R.A. and Williams, T.A. Real-time psychological assessment and evaluation of psychiatric patients. Behavioral Research Methods and Instrumentation 7, 199-200, 1975.
 8. Glueck, B.C. Computers at the Institute of Living. In J.F. Crawford, D.W. Morgan and D. Gianturco (Eds.), Progress in mental health information systems: Computer applications. Cambridge, Mass: Ballinger Publishing Company, 1974.
 9. Laska, E.M. The Multi-state information system. In J.F. Crawford, D.W. Morgan and D. Gianturco (Eds.), Progress in mental health information systems: Computer applications. Cambridge, Mass: Ballinger Publishing Company, 1974.
 10. Sletten, I.W., and Hedlund, J.L. The Missouri automated Standard System of Psychiatry: Current status, special problems and future plans. In J.F. Crawford, D.W. Morgan and D. Gianturco (Eds.), Progress in mental health information systems: Computer applications. Cambridge Mass: Ballinger Publishing Company, 1974.
 11. Spitzer, R.L. and Endicott, J. Can the computer assist physicians in psychiatric diagnosis? American Journal of Psychiatry, 131, 523-530, 1974.
 12. Sletten, I.W. and Hedlund, J.L. The future of computers and actuarial methods in mental health practice. Presented at the International College of Psychosomatic Medicine Symposium IV: Rating Devices and Information Processing in Psychosomatics, Catholic University, Rome, Italy, September 16-20, 1975.

APPENDIX A

OVERVIEW OF ARTIFICIAL INTELLIGENCE RESEARCH

B. G. Buchanan and E. A. Feigenbaum
Stanford University

We give here a brief overview of artificial intelligence (AI) taken from a description of the Stanford Artificial Intelligence Laboratory. The articles following the overview are taken from a preliminary draft of a handbook about AI being written at Stanford under Professor Feigenbaum's supervision. The intent of the articles is to convey some sense of the techniques, problems and successes of AI. Only a few of the most relevant articles are reproduced here.

OVERVIEW

Artificial intelligence is the name given to the study of intellectual processes and how computers can be made to perform them. Some workers in the field believe that it will be possible to program computers to carry out many intellectual process now done by humans. However, almost all agree that we are not very close to this goal and that some fundamental discoveries must be made first. Therefore, work in AI includes trying to analyze intelligent behavior into more basic data structures and processes, experiments to determine if processes proposed to solve some class of problems really work, and attempts to apply what we have found so far to practical problems.

The idea of intelligent machines is very old in fiction, but present work dates from the time stored program electronic computers became available starting in 1949. Any behavior that can be carried out by any mechanical device can be represented in a computer, and getting a particular behavior is "just" a matter of writing a program unless the behavior requires special input and output equipment. It is perhaps reasonable to date AI from A.M. Turing's 1950 paper. Newell, Shaw and Simon started their group in 1954 and the M.I.T. Artificial Intelligence Laboratory was started by McCarthy and Minsky in 1958. [The Stanford AI Lab was started in 1963.]

Board Games. Early work in AI included programs to play games like chess, checkers, kalah and go. The success of these programs was related to the extent that human play of these games makes use of mechanisms we didn't understand well enough to program. If the game requires only well understood mechanisms, computers play better than humans. Kalah is such a game. The best rating obtained in tournament play by a chess program so far is around 1700 which is a good amateur level. The chess programmers hope to do better.

Formal Reasoning. Another early problem domain was theorem proving in logic. This is important for two reasons. First, it provides another area in which our accomplishments in artificial intelligence can be

compared with human intelligence. Again the results obtained depend on what intellectual mechanisms the theorem proving requires, but in general the results have not been as good as with game playing. (This is partly because the mathematical logical systems available were designed for proving metatheorems about logic rather than for proving theorems in logic.)

The second reason why theorem proving is important is that logical languages can be used to express what we wish to tell the computer about the world, and we can try to make it reason from this what it should do to solve the problems we give it. It is quite difficult to express what humans know about the world in the present logical languages or in any other way. Some of what we know is readily expressed in natural language, but much basic information about causality and what may happen when an action is taken is not ever explicitly stated in human speech. This gives rise to the representation problem of determining what is known in general about the world and how to express it in a form that can be used by the computer to solve problems.

Publications. The results of current research in artificial intelligence are published in the journal Artificial Intelligence, and in more general computer science publications such as those of the ACM and the British Computer Society. The ACM has a special interest group on artificial intelligence called SIGART which publishes a newsletter. Every two years there is an international conference on artificial intelligence which publishes a proceedings. The fourth and most recent was held in the U.S.S.R. at Tbilisi in the September 1975 and the proceedings are available.

SUMMARY ARTICLES ON SELECTED TOPICS

The following are selected articles on various aspects of Artificial Intelligence research taken from the current collection of articles in the AI handbook effort. A complete outline of the articles planned can be found in Appendix B. The following articles include discussions of production systems, rote learning, speech understanding, and PLANNER.

PRODUCTION SYSTEMS

GENERAL DESCRIPTION

A PRODUCTION SYSTEM consists of a set of rules (the productions), a data base and an interpreter for the rules. The data base is a collection of symbols. The interpreter tries to match the left hand side of each production to the data base. The interpreter performs the processes on the right hand side of the production if the condition on the left hand side matches some element in the data base. The productions are generally ordered so that if the condition on the left hand side of more than one production matches an element in the data base, the production higher in the order takes priority.

DATA BASE EXAMPLES

The data base of a production system may be simply a set of symbols intended to reflect the state of the world. Some production systems are intended to model a memory mechanism, for example, "short term memory". For these, each element of the data base may represent some piece of knowledge. Examples of systems modeling short term memory are PSG [Newell, 1973] and VIS [Moran, 1973]. Sample elements from the data base for VIS are

```
(HEAR NORTH EAST % END)
(L-2 LINE EAST P-2 P-1)
```

The data bases for knowledge-based experts such as MYCIN [Shortliffe 1975] and DENDRAL [Feigenbaum 1971, Smith 1972] contain facts and assertions about their respective domains of knowledge. For example, the data base in the DENDRAL system contains complex graph structures which represent molecules and molecular fragments. Sample elements from the MYCIN data base are

```
(IDENTITY ORGANISM-1 E.COLI .8)
(SITE CULTURE-2 BLOOD 1.0)
```

A third type of data base is the "token stream approach" in which the data base is a linear stream of tokens accessible only in sequence. An attempt is made to match each production to the beginning of the stream and if a match succeeds, characters in the matched segment may be deleted or modified or new characters may be added. This data base organization was used in LISP70 [Tesler 1973].

In all the production systems described above the data base is the only storage medium for all variables of the system. There is no separate control state information such as a program counter or stack as is used in procedurally-oriented languages. The data base is accessible to every rule in the system and thus serves as a communication channel. The contents of the data base always reflect the current state of the production system.

VARIATIONS OF PRODUCTION SYSTEMS

Production systems have been used in many different programs and programming environments. Many variations of production systems are possible due to differences in the ordering and accessing of rules. The productions are themselves a source of variation in production systems. It is possible to match against the right hand side of the productions instead of the left hand side to obtain a recognizer for symbolic strings. It is also possible to view the left hand side as a goal to be achieved by matching the right hand side of the production. The data base may be a source of variation in production systems as has been discussed above.

ENVIRONMENT

Production systems are particularly appropriate in a domain consisting of a large number of independent states requiring independent actions. The states and actions can be modeled easily using rules, which are also modular in nature. Procedure-oriented systems often find it difficult to update and maintain large numbers of state variables. Production systems are particularly appropriate in this instance. Each production can be viewed as a "demon" ready to be invoked when a particular system state occurs. Production systems are also appropriate where the ability to recognize and react to small variations in the domain is important.

REFERENCES

An excellent reference that discusses in detail many aspects of production systems is a paper by R. Davis and J. King entitled "An Overview of Production Systems", (A.I. Memo 271, Stanford Computer Science Department, November 1975). Other references used in this article are:

Minsky M., Computation Finite and Infinite Machines, Prentice-Hall, 1972.

Newell A., Simon H., Human Problem Solving, Prentice-Hall, 1972.

ROTE LEARNING

BRIEF DESCRIPTION AND HISTORY

Rote learning is a technique which effectively increases the depth of tree searches by recognizing nodes (situations) that have been encountered and evaluated previously. This is done by consulting a file which contains for each node previously encountered a description of that node and the result of the evaluation at that encounter.

This technique was first used by A. L. Samuel in his Checker playing program [See Samuel 1959]. The program used rote learning to accumulate experience over the games it played.

ENVIRONMENT

Rote learning is particularly useful when searching game type trees where the value of a position (node, state) is determined by use of an evaluation function.

TECHNIQUE

Assume that there exists a list of nodes which have been evaluated previously. Associated with each node description is an evaluation. This list will be called the memory file. At the very beginning of the learning process, the file contains nothing in the list. The steps below show how the file is built up by the rote learning process.

The basic steps to rote learning are:

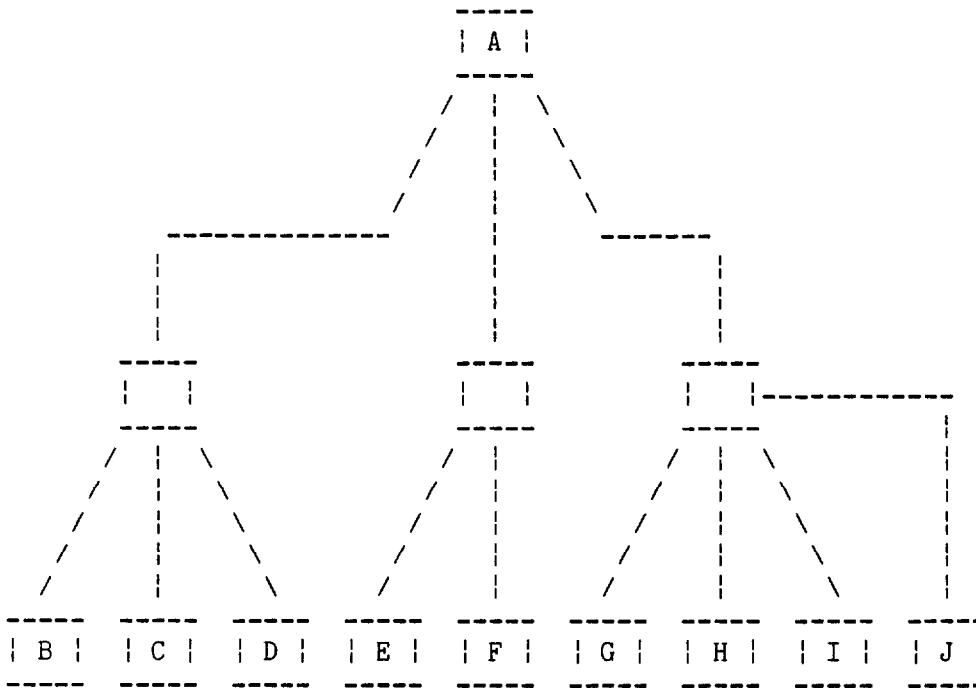
- 1) From the current node which is to be evaluated, form the tree which is to be searched. The form and size of the tree may be governed by a set of heuristics. (For example, expanding the tree fully to a depth of three ply)
- 2) Evaluate the deepest nodes as follows:
 - a) Examine the memory file to see if any of the deepest nodes have been previously evaluated. If so, retrieve from the file the evaluations of these nodes. (Effectively then, these nodes have been evaluated by a further tree search.)
 - b) Any of the deepest nodes not present in the memory file should now be evaluated by the evaluation function.
- 3) Now that all of the deepest nodes have been evaluated back up the tree in the usual min-max fashion to obtain the evaluation for the current node, and to obtain the decision.

- 4) Save a description of the current node and its value in the memory file.

EXAMPLE

Assume that we are playing a game, and that the tree search heuristic is to completely expand the tree to a depth of two ply. Further suppose that we have arrived at a node A, which we wish to evaluate. Assume that this is not our first game, so that the memory file is not empty. We follow the steps of the rote learning procedure as shown:

STEP 1: From node A, we expand the game tree to a depth of two ply, and label the deepest nodes B through J:



STEP 2: Looking in the memory file we discover that nodes B, C, E, F, and H have been previously evaluated, so that we already have their values. Thus we apply the evaluation function only to nodes D, G, I, and J to obtain their values.

STEP 3: Now that the deepest nodes have values, we can go up the tree in the standard fashion, eventually assigning a value to node A, and deciding which branch to take.

STEP 4: Now that we have a value for node A, we place a

description of the node (for instance, the state vector description) and the value of the node in the memory file for possible future use.

BENEFITS OF ROTE LEARNING

As can be seen from the example, since the values of nodes B, C, E, F, and H were retrieved from the memory file, there is in effect a tree search emanating from each of these nodes, and this tree search took place sometime in the past. Thus the depth of the tree search for node A is only 2 ply in some areas, and of greater ply in others. Now if node A itself is ever retrieved from the memory file for evaluation of another node, the depth of the tree for that node is even greater.

LIMITATIONS

- 1) A potential problem with implementing rote search is the storage and retrieval of information from the memory file, especially as the file grows in size. In cases where rote learning is used on a problem of significant size searching the memory file becomes a task which can take the majority of effort. Techniques from other areas of computer science may be used to aid in efficiently maintaining and searching this file. In addition, it is sometimes necessary or desirable to cull the file (delete entries). In this case, heuristics must be devised to determine which entries to keep and which should be purged.
- 2) It is difficult to use rote learning in conjunction with learning schemes which modify the evaluation function (for example, signature tables). The reason for this is that once the evaluation function is changed, in principle every previously evaluated node in the memory file should be re-evaluated, so that the values of newly evaluated nodes and previously evaluated nodes may be meaningfully compared.

COMMENTS

- 1) Samuel found in his Checker playing program that rote learning worked best in the opening and end games. He hypothesized that rote learning functions reasonably well where the results of any specific action are long delayed, or in situations where highly specialized techniques are required (the Checker playing program learned to avoid obvious traps in the end game).
- 2) By slight modification of the procedure described above rote learning can be used with other true search techniques, such as the alpha-beta search, plausibility ordering, or tree-pruning. The fact to realize is that it is not necessary to expand the tree before looking nodes up in the memory file.
- 3) Samuel observed that rote learning can cause nodes which may both lead to winning situations to receive equal weight in decision making, although one of the nodes may lead to a win much more quickly than the

other. Since it is usually desirable to play a shorter game, the depth of the tree search which leads to each node's evaluation should be considered in this case; the node which has a smaller number of plys to the win should be chosen. Thus it may be necessary to store in the memory file the depth of the search for each node stored in the file.

REFERENCES

Samuel, A.L.; "Some Studies in Machine Learning Using the Game of Checkers," IBM Journal 3, 211-229 (1959). Reprinted (with minor additions and corrections) in COMPUTERS AND THOUGHT, edited by Feigenbaum and Feldman, McGraw-Hill, 1963.

SPEECH UNDERSTANDING

Introduction:

The aim of a "speech understanding" system is determination, for spoken utterances, of the intended message in relation to the accomplishment of some task and in spite of indeterminacies and errors in generation, transmission, and reception of the utterance. This is to be distinguished from the aim of a "speech recognition" system, which is provision of an orthographic transcription of the sounds and words corresponding to the acoustic signal. Thus the aim of a speech understanding system does not necessarily include production of an accurate phonetic transcription of the input signal, or an accurate list of the successive words of the input (although it must surely correctly recognize most of them). In other words, if a situation arises in which acoustic processing is unable to resolve the decision between two phonemes or words at a particular point in an utterance, but the overall system is still able to decide the meaning of the sentence, then the sentence is deemed to have been correctly understood.

It seems apparent that a speech recognition system requires a number of different types of processing, each of which corresponds to a different source of information, in order to achieve its aims. It is now well established that knowledge of vocabulary, syntactic, semantic, and pragmatic constraints of a language is required to compensate for errors and uncertainties in the acoustic realization of an utterance.

In summary, a speech understanding system, as presently conceived, will generally fit the following description.

- 1) The system is organized into a number of levels, starting with the acoustic and working up to the syntactic and semantic.
- 2) Action is generally from the lower levels upward, utilizing programs that incorporate knowledge of each particular level.
- 3) Task limitations are used at several levels to help make selections.
- 4) The higher levels are sometimes used in a feedback mode at lower levels to help make selections.

History:

Speech recognition research has yielded significant results in the case of isolated words (accuracy greater than 95%). The primary emphasis has been on acoustic processing and classical pattern re-cognition and matching techniques. Straight-forward extrapolation of these techniques to continuous speech recognition, however, has not proved successful. It is felt that a major reason for the difficulties encountered is that the information used by humans in understanding speech is not completely contained in the acoustic representation of the speech signal. Experiments by Klatt and Stevens (1972) in the area of spectrogram reading showed that

the performance obtained by human experts for phonetic segmentation and labelling without conscious appeal to syntactic, semantic, and vocabulary constraints was: approximately 75% correctly labelled, 15% mislabelled, and 10% missed. When these other sources of knowledge were used, the success rate for word identification rose to 96%. These results have greatly influenced recent research in speech understanding.

Possible Applications:

Speech would be an appropriate input channel to a computer in many situations. The average output data rate is higher for speech than for writing or typing. Use of the speech channel does not tie up other effectors, such as hands, eyes, feet, or ears. It can therefore be used while in motion or in parallel with other channels. Speech is also a preferred channel for spontaneous communication of the type that is found in an interactive environment.

Long range applications are readily listed. They might include for example, automatic dictation systems, voice-response order takers, or in the computer area, a voice operated graphics terminal.

In the shorter term, several tasks have been suggested as possible vehicles for research in speech understanding (Newell et al 1973). They are:

- 1) Querying a Data Management System
- 2) Data Acquisition of Formatted Information
(voice-key-punch)
- 3) Querying the Operational Status of a Computer
- 4) Consulting on the Operation of a Computer (i.e., a voice-operated HELP)

Unsolved Problems:

The following is a brief discussion of unsolved problems in speech understanding following Newell(1973) and roughly ordered in terms of system level (i.e. from acoustic at the lowest to semantic at the highest).

The essential problem of continuous speech at the acoustic level is phoneme-level identification and not necessarily segmentation between words. There is, however, a significant amount known about acoustic-phonetic and phonological rules which has yet to be fully exploited in production systems. The difficulty of adapting to multiple speakers of different sexes and with different dialects also remains a problem, although it is hoped that proper normalization of acoustic-phonetic and phonological rules will make them speaker-invariant. Two other acoustic-related problems are environmental noise and possible distortions caused by the communications channel (e.g. the telephone channel).